vmware SOLUTION READINESS

Citrix XenApp on VMware
Best Practices Guide

© 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at http://www.vmware.com/download/patents.html.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

© 2011 VMware, Inc. All rights reserved.

Contents

Contents 3
1. Overview 4
1.1 Purpose 4
1.2 Target Audience 4
1.3 Scope 4
2. Citrix XenApp Architecture on vSphere 5
3. VMware ESX Host Best Practices for Citrix XenApp 6
3.1 CPU Configuration Guidelines 6
3.2 Memory Configuration Guidelines 8
3.3 Storage Guidelines 9
3.4 Networking Guidelines 12
4. Citrix XenApp on vSphere Best Practices 14
4.1 Virtual Hardware 14
4.2 Guest Operating System 15
4.3 Citrix XenApp Best Practices 17
5. Monitoring Performance 18
5.1 Monitoring vSphere 18
5.2 Monitoring Citrix 21
6. vSphere Enhancements for Deployment and Operations 22
6.1 Rapid Provisioning and Deployment 22
6.2 Efficient Change Management 22
6.3 Enhanced Availability with VMware HA 23
6.4 Resource Optimization with VMware DRS and vMotion 24
6.5 Security with VMware vShield Endpoint 25

© 2011 VMware, Inc. All rights reserved.

1. Overview

Desktop application delivery and management can be tedious and time-consuming. Many organizations have chosen to leverage application virtualization and take a software as a service (SaaS) approach to desktop applications. By deploying software such as Citrix XenApp IT organizations are able to centralize the management, administration and delivery of popular Windows-based applications. This model of centralized administration and delivery of services is not new to VMware customers who have for years used virtualization technology to consolidate server workloads.

This guide provides information about deploying Citrix XenApp in a virtualized environment powered by VMware vSphere™. Key considerations are reviewed for architecture, performance, high availability, and design and sizing of virtualized workloads, many of which are based on current customer deployments of XenApp application servers on VMware. This guide is intended to help IT Administrators and Architects successfully deliver virtualized applications using Citrix XenApp on VMware vSphere.

1.1 Purpose

This guide provides best practice guidelines for deploying Citrix XenApp on VMware. The recommendations are not specific to any particular set of hardware or to the size and scope of any particular XenApp implementation. The design examples and considerations provide guidance only and do not represent strict design requirements. Flexibility in Citrix XenApp deployments combined with the flexibility of deploying on vSphere allows for a wide variety of valid configurations.

1.2 Target Audience

A basic knowledge and understanding of VMware vSphere and Citrix XenApp is assumed.

- Architectural Staff can use this document to gain an understanding of how the system will work as a whole as they design and implement various components.
- Engineers and administrators can use this document as a catalog of technical capabilities.

1.3 Scope

This best practices guide focuses on the following topics:

- Citrix XenApp Architecture on vSphere – Provides background on Citrix XenApp architecture and the rationale for deploying on VMware vSphere.
- VMware ESX™ Host Best Practices for Citrix XenApp –Provides proven VMware best practices for vSphere hosts running XenApp workloads. Includes guidance in the areas of CPU, memory, storage, and networking.
- Citrix XenApp on vSphere Best Practices – Deploying Citrix XenApp on vSphere requires that proven best practices for the XenApp application continue to be followed. The focus in this section is on configuring virtual machines for XenApp.
- Monitoring Performance – When migrating XenApp to a vSphere infrastructure, maintaining performance levels that are equal or better than those achieved in physical deployments is essential. Monitoring before and after a migration helps validate whether the migration was a success, and can also help establish a baseline understanding of the performance characteristics. This section takes a look at the vSphere tools available to help monitor the vSphere environment.
- vSphere Enhancements for Deployment and Operations – Provides a brief look at vSphere features and add-ons that can enhance the deployment and management of XenApp.

© 2011 VMware, Inc. All rights reserved.

2. Citrix XenApp Architecture on vSphere

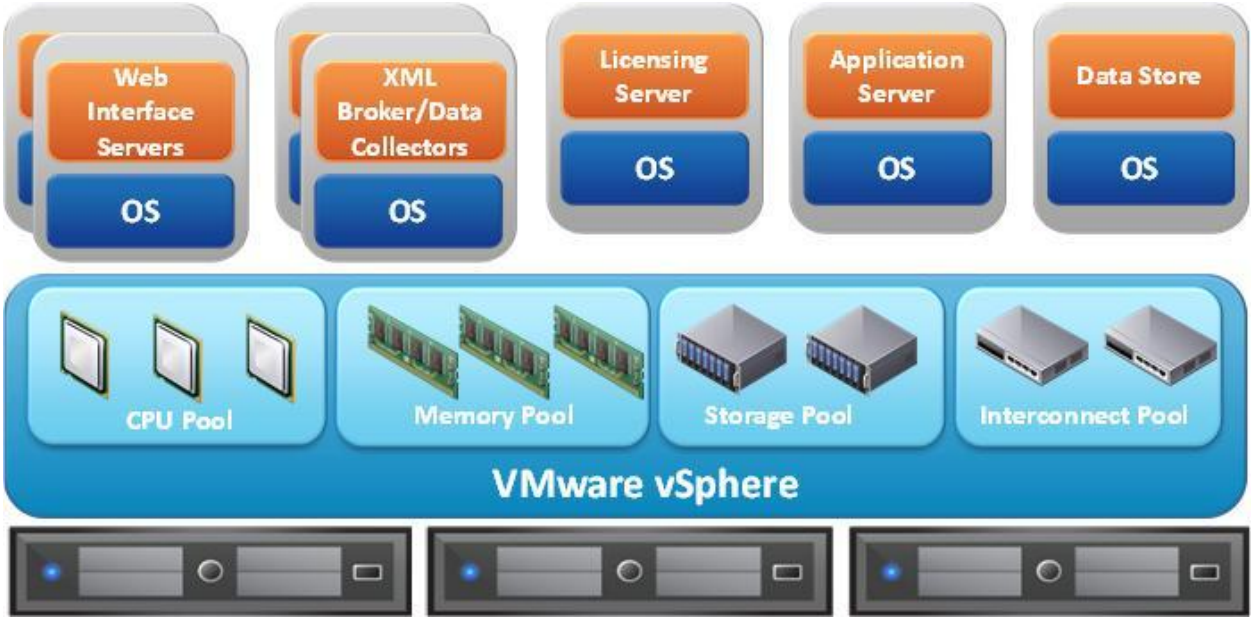
As with any multi-tiered application, XenApp can be spread across many individual servers, or consolidated into a few. Consolidation is typical in physical environments where there is either a lack of servers to adequately segregate the components, or where the environment is too small to justify deploying multiple physical servers. Consolidation has also been driven by higher capacity server hardware and the fact that most applications cannot fully take advantage of the resources provided. Scaling the components out is typically a best practice from a security and availability perspective; however, this requires a significant initial investment when deploying in a non-virtualized environment.

The Citrix XenApp architecture consists of, at a minimum, the following components:

- Licensing Server – Required for all XenApp deployments.
- Data Store – Database where information is stored, such as configuration information for published applications, users, printers, and servers.
- Data Collector – Maintains information such as license usage, session status, server loads, users connected, and published applications.
- Web Interface – Required when users access applications using the online plug-in or a web browser.
- XML Service and Broker – Intermediary between the web interface and the XenApp servers.
- Application Servers – Hosts the published applications to which users connect.

Deploying on vSphere provides the ability to isolate each role into its own virtual machine and operating system. This allows you the greatest flexibility and deployment options. By keeping the components separate you can be sure that if a virtual machine has a problem only the specific component will be affected. Having the ability to deploy multiple virtual machines can also help increase the availability of the component. Figure shows the required XenApp components deployed on vSphere.

Figure 1. XenApp on vSphere



© 2011 VMware, Inc. All rights reserved.

3. VMware ESX Host Best Practices for Citrix XenApp

Before XenApp virtual machines can be built and applications deployed a solid virtual infrastructure must be in place. The following are key vSphere best practices for properly designing and deploying ESX hosts to run XenApp.

3.1 CPU Configuration Guidelines

As processor technology continues to evolve it's becoming more difficult to make efficient use of new multi-core servers. Many organizations are turning to virtualization for exactly this reason. Understanding the capabilities of newer processor technologies and how they can be used more efficiently in a virtualized environment is key to the foundation of a well-designed vSphere environment.

3.1.1 Use Latest CPU Technology

When possible, deploy vSphere on CPUs that support second generation CPU virtualization. This includes processors from both AMD and Intel. Second generation CPU virtualization not only supports hardware assisted CPU virtualization (Intel VT-x and AMD AMD-V), but also introduces the hardware assisted memory management unit (MMU). AMD refers to this technology as rapid virtualization indexing (RVI) or nested page tables (NPT). On Intel processors this feature is called extended page tables (EPT). Hardware assisted MMU eliminates the overhead required for ESX to maintain mappings in software of guest virtual memory to host physical memory addresses.

3.1.2 Use vSphere 4.0 or Later

With the release of vSphere 4.0, VMware introduced improvements that take advantage of new processor technology. These improvements include:

- *Relaxed co-scheduling* has been further improved in vSphere 4.0. This allows for greater efficiency when running multiple SMP virtual machines. Relaxed co-scheduling requires only skewed vCPUs be co-started, as opposed to *strict co-scheduling,* which requires all vCPUs to be co-started.
- Fine-grained locking is used to reduce scheduling overheads in cases where frequent scheduling decisions are required.
- The new scheduler is aware of processor cache topology and takes into account the processor cache architecture to optimize CPU usage.

3.1.3 CPU Over-Commitment

A useful side effect of virtualization is the abstraction of the underlying hardware. Though this is one of the great benefits of virtualization, it can cause significant performance issues if overcommitment of resources is not well understood. As a best practice, the initial deployment of virtual processors on a host should not exceed the number of physical processors. For example, a host with 16 physical cores could initially host four 4-vCPU virtual machines or eight 2-vCPU virtual machines. This is not to say that the host cannot be overcommitted; rather, a baseline should be established prior to overcommitting to be able to fully understand the workload that is being applied to the virtual machines. Once a baseline is established workloads can be added until a balance of overcommitment and acceptable application latency has been achieved.

© 2011 VMware, Inc. All rights reserved.

3.1.4 Enable Hyper-Threading

If the processor technology supports hyper-threading, also called *symmetric multithreading* or SMT, the VMware best practice is to enable this feature within the BIOS. Hyper-threading allows a single processor to behave like two logical processors, allowing two independent threads to run simultaneously. It's important to understand that two logical processors do not double the performance, but they do provide an increase in performance by better utilizing idle resources leading to greater throughput.

3.1.5 Enable NUMA

Non-Uniform Memory Access (NUMA) compatible systems contain multiple nodes that consist of a set of processors and memory. The access to memory in the same node is local, while access to the other node is remote. Remote access can take longer because it involves a multi-hop operation. In NUMA aware applications attempt to keep threads local to improve performance.

ESX provides load-balancing on NUMA systems. To achieve the best performance it is recommended that NUMA is enabled on compatible systems. On a NUMA-enabled ESX host, virtual machines are assigned a home node from which the virtual machine's memory is allocated. Because it is rare for a virtual machine to migrate away from the home node, memory access is mostly kept local. See Section 3 for more information about how to best utilize the NUMA capabilities of ESX.

3.1.6 Halting Idle Millisecond Penalty Parameter (HIMP)

ESX's CPU scheduler uses resource settings such as reservations, limits and shares to make sure that a virtual machine receives the CPU time that it is entitled to. If a virtual machine falls behind the scheduler may attempt to schedule it more frequently to help it catch up. When hyper-threading is enabled more frequent scheduling may not guarantee that the virtual machine will catch up. This is due to the fact that on a single physical core one logical processor core is affected by the activity of the other logical processor core. To guarantee that the virtual machine can catch up, ESX may opt not to schedule virtual machines on the other logical core, thus leaving it idle.

Not scheduling on a logical processor may have a measureable impact on overall throughput of the system. This can occur when the following conditions are true:

- Modern processors with hyper-threading (for example, Nehalem)
- More than 50% CPU utilization
- Number of vCPUs = number of pCPUs +/- 25%
- CPU usage patterns show frequent spikes

VMware knowledge base article KB 1020233 provides guidance relating to the HIMP parameter. Testing may prove that changes to this parameter provide performance gains within your environment. If so, it is strongly recommended to refer back to the KB article on a regular basis as guidance may change with later releases of vSphere.

© 2011 VMware, Inc. All rights reserved.

3.2 Memory Configuration Guidelines

This section covers concepts used for memory management in vSphere and techniques used to make sure that virtual machines get the memory they require for their respective workload.

3.2.1 ESX Memory Management Concepts

vSphere virtualizes guest physical memory by adding an extra level of address translation. Hardware-assisted virtualization technologies make it possible to provide this additional translation with little or no overhead. Managing memory in the hypervisor enables the following:

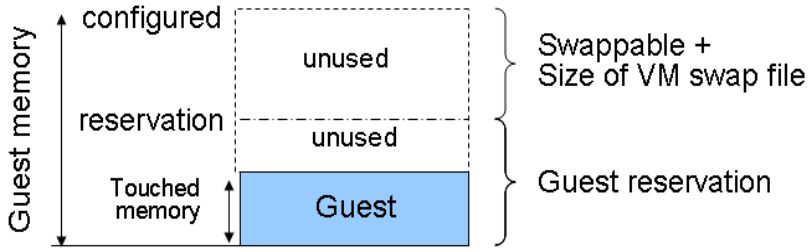
- Memory sharing across virtual machines that have similar data (same guest operating systems).
- Memory overcommitment, which means allocating more memory to virtual machines than is physically available on the ESX host. Overcommitment is not necessarily a bad thing. Many customers can achieve high levels of consolidation and efficiency using it. However, overcommitment should be carefully monitored to avoid negative performance impact.
- A memory balloon technique, whereby virtual machines that do not need all the memory they have been allocated give memory to virtual machines that require additional allocated memory.

For more details about vSphere memory management concepts, see *Understanding Memory Resource Management in VMware ESX 4.1*.

3.2.2 Virtual Memory Concepts

Figure 2 illustrates the use of memory settings parameters in the virtual machine.

Figure 2. Virtual Machine Memory Settings



The vSphere memory settings for a virtual machine include the following parameters:

- Configured memory – Memory size of virtual machine assigned at creation.
- Touched memory – Memory actually used by the virtual machine. vSphere only allocates guest operating system memory on demand.
- Swappable – Virtual machine memory that can be reclaimed by the balloon driver or by vSphere swapping. Ballooning occurs before vSphere swapping. If this memory is in use by the virtual machine (touched and in use), the balloon driver causes the guest operating system to swap. vSphere 4.1 introduces memory compression (discussed in the following section).
- If the balloon driver is unable to reclaim memory quickly enough or is disabled or not installed, vSphere forcibly reclaims memory from the virtual machine using the VMkernel swap file.

© 2011 VMware, Inc. All rights reserved.

3.2.3 Memory Best Practices

The following are best practices for memory in ESX.

- Account for memory overhead – Virtual machines require memory beyond the amount allocated, and this memory overhead is per-virtual machine. Memory overhead includes space reserved for virtual machine devices, such as SVGA frame buffers and internal data structures. The amount of overhead required depends on the number of vCPUs, configured memory, and whether the guest operating system is 32- or 64-bit.
- ”Right-size” memory allocations – Over-allocating memory to virtual machines can not only unnecessarily waste memory, but increases the amount of memory overhead allocated to the virtual machine, thus reducing the overall memory available for other virtual machines.
- Use vSphere 4.1 – New features introduced in vSphere 4.1 can enhance memory performance. One technique, *memory compression*, was introduced to provide an additional technique to combat performance degradation due to memory overcommitment. Before resorting to ESX host swapping, pages in memory attempt to be compressed and stored in the virtual machine’s compression cache. The decompression latency is much smaller than the swap-in latency, so this can greatly improve system performance in overcommitment scenarios.
- Use Transparent Page Sharing – Redundant pages in memory can be reclaimed by the use of transparent page sharing. At run-time, virtual machines may have identical sets of memory content (for example several virtual machines running the same guest operating system). Transparent page sharing lets virtual machines share these pages, thus increasing the memory available on an ESX host.

3.3 Storage Guidelines

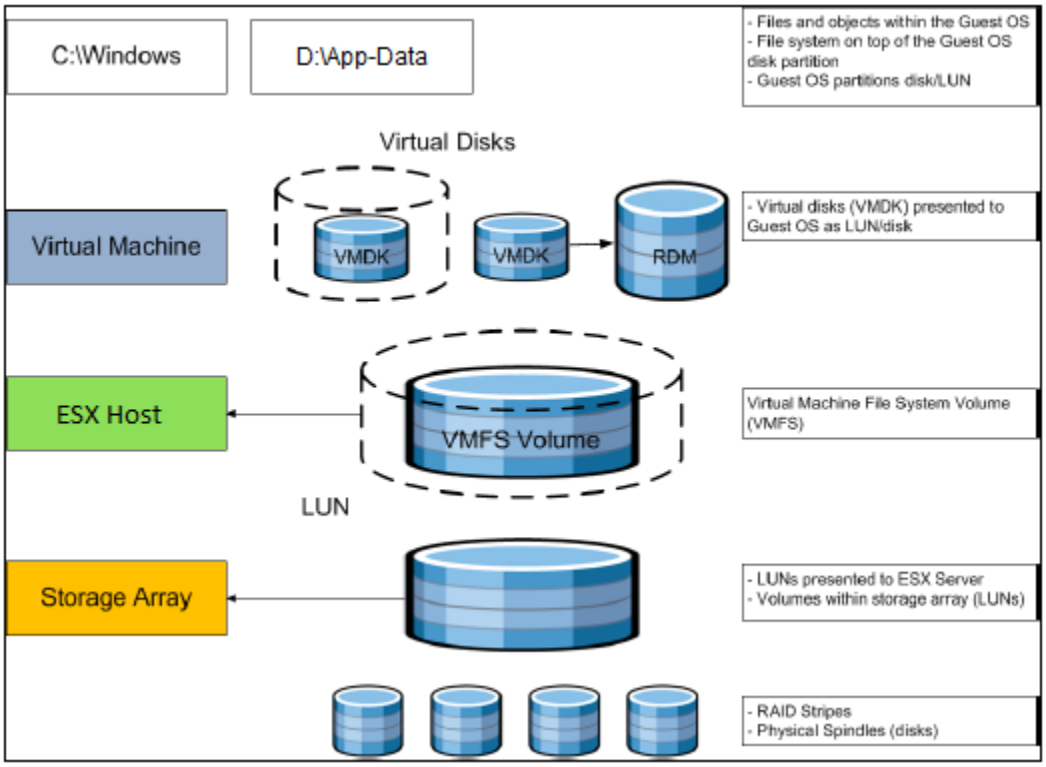
vSphere provides many features that take advantage of commonly used storage technologies such as storage area networks and storage replication. Features such as VMware vMotion™, VMware HA and DRS use these storage technologies to provide high-availability, resource balancing and uninterrupted workload migration.

3.3.1 Storage Virtualization Concepts

As illustrated in Figure 3, VMware storage virtualization can be categorized into three layers of storage technology. The Storage Array is the bottom layer, consisting of physical disks presented as logical disks (storage array volumes or LUNs) to the layer above, with the virtual environment occupied by vSphere. Storage array LUNs that are formatted as VMFS volumes that provide storage for virtual disks. Virtual machines consist of virtual disks that are presented to the guest operating system as SCSI attached disks that can be partitioned and used in file systems.

© 2011 VMware, Inc. All rights reserved.

Figure 3. VMware Storage Virtualization Stack



3.3.1.1. VMFS File System

The VMFS file system was created by VMware to allow multiple vSphere hosts to read and write to the same storage concurrently. VMFS is a clustered file system that allows you to simplify virtual machine provisioning and administration by consolidating virtual machines into smaller units of storage. Unique virtualization-based capabilities provided by the VMFS file system include live migration using vMotion, and increased availability using VMware HA.

Virtual machines are stored on VMFS file systems as a unique set of encapsulated files, including configuration files and virtual disks (VMDK files). VMFS is supported on both iSCSI and Fibre Channel attached storage.

3.3.1.2. Raw Device Mapping

For instances where isolation or direct access to the underlying storage subsystem is required a raw device mapping can be used in place of virtual disks. Raw device mappings use a mapping file that is located on a VMFS volume to point to a physical LUN. The physical LUN is accessible to the virtual machine in its raw form and must be formatted from within the virtual machine. Unlike VMFS, a raw device mapping is typically only assigned to a single virtual machine. However, RDMs can be shared; for example, in a Microsoft Cluster configuration where multiple nodes use SCSI reservations to handle arbitration. RDMs cannot provide all of the features available with VMFS and should be limited to use only when technically required.

© 2011 VMware, Inc. All rights reserved.

3.3.2 Storage Protocol Capabilities

VMware vSphere provides vSphere and storage administrators with the flexibility to use the storage protocol that meets the requirements of the business. This can be a single protocol datacenter wide, such as iSCSI, or multiple protocols for tiered scenarios such as using Fibre Channel for high-throughput storage pools and NFS for high-capacity storage pools.

For XenApp on vSphere there is no single option that is considered superior to another. It is recommended that this decision be made based on your established storage management practices within the virtualized environment.

Refer to the VMware whitepaper *Comparison of Storage Protocol Performance in VMware vSphere 4* for details.

3.3.3 Storage Best Practices

The following are vSphere storage best practices.

- Host multi-pathing – Having a redundant set of paths to the storage area network is critical to protecting the availability of your environment. This redundancy can be in the form of dual host-bus adapters connected to separate fabric switches, or a set of teamed network interface cards for iSCSI and NFS.
- Partition alignment – Partition misalignment can lead to severe performance degradation due to IO operations having to cross track boundaries. Partition alignment is important both at the VMFS file system level as well as within the guest operating system. Use the vSphere Client to create aligned partitions.
- Use shared storage – In a vSphere environment many of the features that provide the flexibility in management and operational agility come from the use of shared storage. Features such as VMware HA, DRS, and vMotion take advantage of the ability to migrate workloads from one host to another host while reducing or eliminating the downtime required to do so.
- Calculate your total virtual machine size requirements – Each virtual machine will require more space than just that used by its virtual disks. Consider a virtual machine with a 20GB OS virtual disk and 16GB of memory allocated. This virtual machine will require 20GB for the virtual disk, 16GB for the virtual machine swap file (size of allocated memory), and 100MB for log files, (*total virtual disk size* + *configured memory* + *100MB)* or 36.1GB total.
- Understand IO Requirements – Under-provisioned storage can significantly slow responsiveness and performance for end users. Operations such as logon activity at the beginning of the day or antivirus signature updates throughout the day can quickly bring under-provisioned storage to light. If an existing XenApp environment is in place, use tools such as Windows Performance Monitor (physical environment) or VMware esxtop (VMware environment) to monitor disk IO throughput. This provides a baseline that can be used to properly size the storage infrastructure. If no XenApp environment is in place, perform load and scalability tests using tools such as Citrix EdgeSite for Load Testing to accurately size the storage infrastructure.

© 2011 VMware, Inc. All rights reserved.

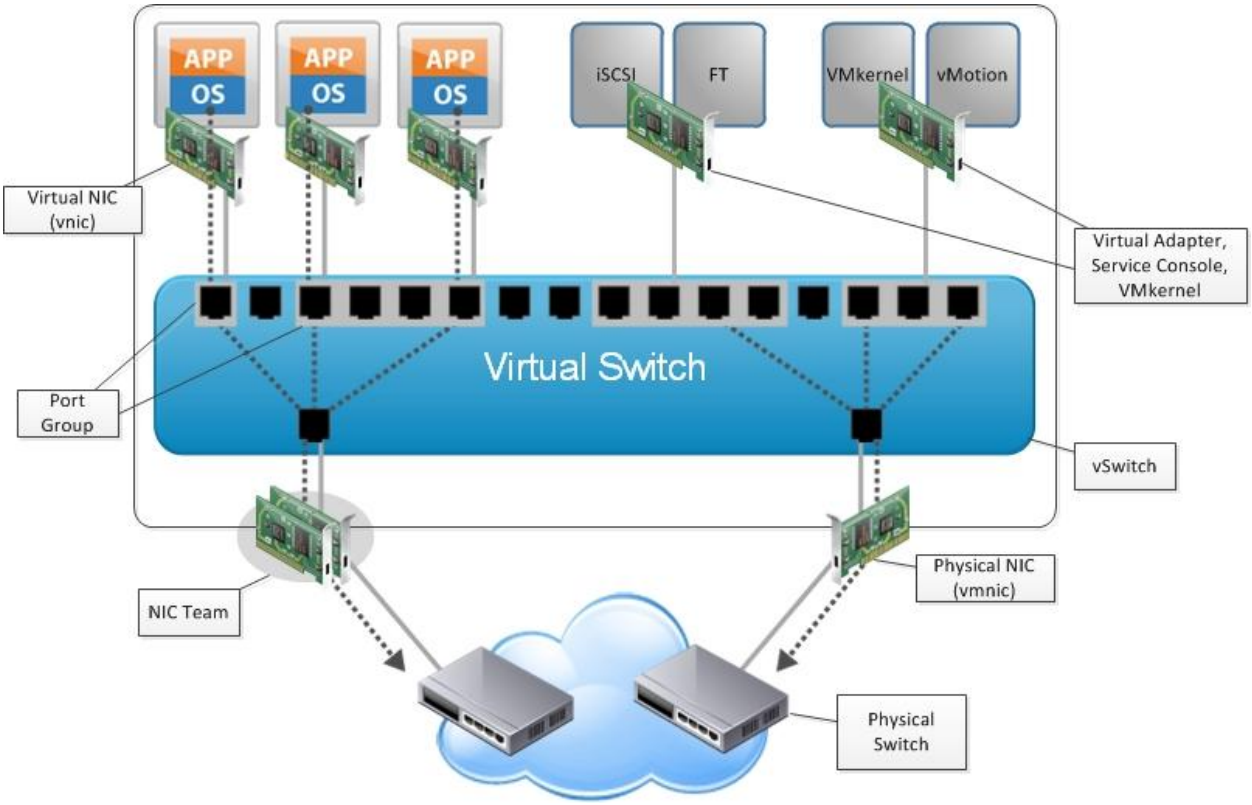
3.4 Networking Guidelines

Networking in the virtual world follows the same concepts as in the physical world, but these concepts are applied in software instead of using physical cables and switches. Many of the best practices that apply in the physical world continue to apply in the virtual world, but there are additional considerations for traffic segmentation, availability, and making sure the throughput required by services hosted on a single server can be fairly distributed.

3.4.1 Virtual Networking Concepts

Figure 4 provides a visual overview of the components that make up the virtual network.

Figure 4. Virtual Networking in vSphere



As shown in the figure, the following components make up the virtual network:

- Physical switch – vSphere host-facing edge of the local area network.
- Physical network interface (vmnic) – Provides connectivity between the ESX host and the local area network.
- vSwitch – The virtual switch is created in software and provides connectivity between virtual machines. Virtual switches must uplink to a physical NIC (also known as *vmnic*) to provide virtual machines with connectivity to the LAN. Otherwise, virtual machine traffic is contained within the virtual switch.
- Port group – Used to create a logical boundary within a virtual switch. This boundary can provide VLAN segmentation when 802.1q trunking is passed from the physical switch, or can create a boundary for policy settings.
- Virtual NIC (vNIC) – Provides connectivity between the virtual machine and the virtual switch.

© 2011 VMware, Inc. All rights reserved.

- Vmkernel (vmknic) – Interface for hypervisor functions such as connectivity for NFS, iSCSI, vMotion and FT logging.
- Service Console (vswif) – Interface for the service console present in ESX Classic. Not present in ESXi.
- Virtual Adapter – Provides Management, vMotion, and FT Logging when connected to a vNetwork Distributed Switch.
- NIC Team – Group of physical NICs connected to the same physical/logical network(s) providing redundancy.

3.4.2 Virtual Networking Best Practices

The following are vSphere networking best practices:

- Separate virtual machine and infrastructure traffic – Keep virtual machine and vmkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation.
- Use NIC Teaming – Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches. Teaming provides redundancy against NIC failure and, if connected to separate physical switches, against switch failures. NIC teaming does not necessarily provide higher throughput.
- Enable PortFast on ESX host uplinks – Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. To prevent this situation, switch ports connected to ESX hosts should be set to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESX hosts from affecting the STP topology. Loops are not possible in virtual switches.
- Converge Network and Storage IO with 10Gb Ethernet – When possible consolidating storage and network traffic can provide simplified cabling and management over having to maintain separate switching infrastructures.

© 2011 VMware, Inc. All rights reserved.

4. Citrix XenApp on vSphere Best Practices

Virtual machines hosting desktop applications require slightly more attention than those hosting server applications because there is an inherent assumption that they are running directly on the hardware with which the end user is interacting. A server application builds abstraction into the communication engine that is used to relay information back and forth from the client to the server. Because of this lack of abstraction in the desktop world there are special considerations to be taken into account when building virtual machines to host desktop applications, and these can be mitigated by following best practices for deploying virtual machines to be used as hosts for XenApp.

4.1 Virtual Hardware

Virtual machines allow for flexibility in configuration both at the virtual hardware level as well as in the software configuration. By pre-configuring a *golden template* for a specific application you can quickly deploy new instances and minimize the need for custom configuration. When building a golden template for XenApp, keep in mind the following best practices:

- Building blocks – Test to find the optimal virtual machine configuration for your environment based on your workload and user profile. Use the following tips to determine a configuration that can be considered as the ideal building block for your environment. Use this building block to increase capacity as required in a manageable and predictable fashion. A building block consists of a standard set of virtual hardware, guest OS configurations, and defined number of supported users.
- Build new virtual machines – Avoid converting existing physical machines (P2V). Physical servers require hardware drivers and typically have vendor-specific utilities for monitoring and configuration. These can all add overhead when converted to virtual machines. Additionally, any problems found in the physical deployment are brought over during the conversion.
- Use 1-2 vCPUs – Field engagements and internal testing has found that scaling out (using more, smaller virtual machines) provides improved performance over deploying larger virtual machines in a XenApp environment. Smaller virtual machines also provide easier units to place for VMware Distributed Resource Scheduling (DRS), discussed in Section 5.
- Allocate 4-8GB of memory per virtual machine – Similar to vCPU allocation, the goal is to take a scale-out approach to deploying XenApp on vSphere. The exact amount of memory required is highly dependent on the applications being deployed within the XenApp virtual machine. Start off with a smaller amount of memory and monitor to avoid over-provisioning the virtual machine and wasting memory.
- Size the virtual machine to fit within the NUMA node – ESX is NUMA aware and can help virtual machines maintain memory and CPU locality. The key to leveraging this capability is to keep the virtual machine within the size of the NUMA node. For example, if the host is configured with 2 x quad-core processors and 16GB of memory, the ideal virtual machine size is 4 vCPUs or less and 8GB or less.
- Disable unnecessary virtual hardware – Additional virtual hardware takes up resources. Keeping the virtual machines trim can help achieve that last bit of performance capability. Disable virtual devices such as CD-ROM and floppy drives.
- Use VMXNET3 – This paravirtualized VMware network interface driver provides the best performance while reducing host processing required for networking IO compared to the e1000 driver.

© 2011 VMware, Inc. All rights reserved.

- Use LSI Logic vSCSI adapter – For low IO workloads the LSI Logic virtual SCSI drivers offer the best performance. VMware provides a paravirtualized SCSI adapter, but the general recommendation is to use this driver for high IO workloads (>2000 IOPS). Due to its wide use the LSI Logic driver helps maintain consistency between virtual machine configurations.
- Use thick disks – Thin-provisioned disks may be used, but strict control and monitoring is required to make sure adequate performance is maintained and storage is not completely consumed. If operational tactics are in place to mitigate the risk of performance and storage depletion thin disks are a viable option. Otherwise, the general recommendation is to deploy thick disks.
- Consider whether you should use a separate virtual disk for the OS page file – In the physical world this has long been the recommendation, mostly due to the lack of addressable space in 32-bit operating systems and available direct-attached storage that could be used to house the page file on separate physical disks. 64-bit operating systems have mostly addressed the lack of memory, but page files are still required for kernel dumps and to provide space on disk to move lower priority workloads when high priority processes demand it. In a virtual environment it may not be feasible to place page files on dedicated storage. Additionally, the increase in available memory has reduced the requirement for large page files and the overall paging that occurs. The recommendation is generally to size the OS virtual disk appropriately to accommodate the page file. After the system is under steady load, preferably during the testing phase, monitor to validate the design and make adjustments if necessary.
- Consider setting memory reservation – With applications providing direct user interaction having physical memory available is important to the end-user experience. Reservations have additional benefits beyond guaranteeing physical memory for the virtual machine. Setting a reservation sized equal to the size of memory allocated eliminates the virtual machine swap file. This can reduce storage capacity requirements, and if storage replication is in place it reduces the data that is replicated. Reservations can also be used as a way of limiting overcommitment as a virtual machine cannot be powered on if memory reservation requirements cannot be met. Reservations can add management overhead, so measure and evaluate the pros and cons before standardizing on this practice.

4.2 Guest Operating System

The guest operating system configuration must undergo the same scrutiny as the virtual machine and vSphere infrastructure to optimize the environment to run XenApp. The following best practices have been compiled from successful XenApp on vSphere implementations.

- Use 64 bit OS – 32-bit operating systems are limited to 2GB of user addressable memory. Using techniques to increase this space (`/3gb` and `/pae` in `boot.ini`) can result in kernel memory starvation and overall performance degradation. Using a x64-based OS eliminates the memory limitations of 32-bit operating systems. Application performance should be monitored when deploying 32-bit applications on x64 systems as there is slight overhead due to the emulation required.

  Note If required, virtual machines running 32-bit operating systems can be deployed on the same ESX host as x64-based virtual machines without requiring new or separate hardware or virtual infrastructure.

- Build Windows 2003 templates as multi-vCPU virtual machines – If Windows 2003 is used as the guest operating system build the golden template as a multi-vCPU virtual machine. Templates built as uniprocessor virtual machines only recognize a single vCPU. Multi-vCPU virtual machines deployed using the uniprocessor template will also only recognize a single vCPU until the multi-processor HAL is installed using Device Manager. Building the golden template as a multi-vCPU virtual machine causes all assigned virtual processors to be recognized.

  Note Windows 2008 dynamically adjusts the HAL to support the underlying hardware.

© 2011 VMware, Inc. All rights reserved.

- Align guest OS partitions – Formatted virtual disks can suffer from performance degradation if misaligned. Windows 2008 automatically aligns all partitions to a 1MB boundary. If using Windows 2003 use the command line tool `diskpart` to align the partition. Refer to the following code snippet to align a Windows 2003-based partition.

```
C:\>diskpart
Microsoft DiskPart version 5.2.3790.1830
Copyright (C) 1999-2001 Microsoft Corporation.
On computer: ha-xenapp-1
DISKPART> select disk 1
Disk 1 is now the selected disk.
DISKPART> create partition primary align=64
DiskPart succeeded in creating the specified partition.
```

- Install VMware Tools – VMware provides drivers and guest OS optimizations through the VMware Tools suite. These include video resolution improvements, memory optimization, and visibility into the guest OS from ESX and VMware vCenter™ Server. VMware tools is required to take advantage of the virtual machine monitoring functionality of VMware HA (see Section 5 for additional information).
- Set fixed page file size – Follow Microsoft recommendations for sizing the page file, typically 1.5 times the size of configured memory. Configure the minimum and maximum page file sizes as recommended. This disables dynamic resizing of the page file and helps reduce fragmentation due to resizing.

  Note In x64-based operating systems configured with a large memory size it may not be required to configure such a large page file. Follow Microsoft KB article KB889654 to determine the optimal page file size for the workload.

- Use default balloon driver settings – The balloon driver is installed as part of the VMware Tools suite and can be used by ESX if physical memory comes under contention. Performance tests have shown that the balloon driver allows ESX to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESX to use host-swapping to make up for the lack of available physical memory. Host-swapping can severely impact performance.
- Use transparent page sharing defaults – Redundant copies of memory pages, such as those used by similar operating system images, are eliminated by the use of transparent page sharing (TPS). This allows memory to be freed up for use by virtual machines. TPS is invoked as memory comes under contention and, as a result, a host may show few or zero shared pages if there is no overcommitment. TPS settings are automatically set by vSphere during the creation of the virtual machine based on the chosen operating system. As a best practice leave TPS settings at the default.
- Disable Shared Folders – If VMware Tools is installed using the **Complete** option the Shared Folder feature is installed. Shared folders are used in hosted VMware products such as VMware Workstation and VMware Fusion to allow the mapping of folders available within the host OS to the guest OS. Because the feature is not supported on ESX, disabling it (if it was installed) has no impact. To avoid having the shared folder feature installed use the **Typical** installation option when installing VMware Tools or follow the steps in VMware KB article **KB1317**.
- Antivirus – Consult with your antivirus vendor to obtain the best practices for configuring their respective AV products for XenApp. Additionally, VMware can provide out-of-band antivirus scanning using VMware vShield™ Endpoint (discussed in Section 5).
- Disable unnecessary services – Out of the box Windows installations enable many services that may go unused, yet they consume compute resources and can increase the attack surface. Review the services that have a startup type of **Automatic** and evaluate whether they can be disabled in your environment.

© 2011 VMware, Inc. All rights reserved.

4.3 Citrix XenApp Best Practices

The following, though not a comprehensive list, provides best practices that should be followed in any Citrix XenApp environment. These practices are often overlooked, but can provide great benefits in the environment.

- SMB Tuning **–** SMB 1.0 environments suffer from inherent limitations that can result in poor performance for end users. These include delays when opening documents located in a Windows file share, programs that appear to stop responding, and high CPU utilization. These problems can be mitigated by performing SMB client and server tuning. Refer to Microsoft KB article KB324446 for the recommended registry modifications that should be made to all XenApp servers and Windows file servers.
- Test for scale **–** Scalability testing of a new environment is often overlooked due to marketing claims, or a published scalability test that may have no similarities to the newly created environment. Proper testing using characteristics that closely resemble the expected workload is critical to obtaining useable data about the user density that can be achieved. Though some customers choose to develop their own in-house testing mechanisms, the expertise to develop such tools may not be available in all organizations. Consider using tools such as Citrix EdgeSight for Load Testing which can be used to simulate user load in your XenApp environment.

© 2011 VMware, Inc. All rights reserved.

5. Monitoring Performance

5.1 Monitoring vSphere

Proactive monitoring is essential to understand the performance characteristics of a workload and to provide a useable baseline. The baseline provides data that can be used when you begin looking at alerting and performance troubleshooting. Applications and operating systems provide different methods for capturing this data, but it is useful to monitor the base infrastructure and understanding application characteristics from that perspective.

5.1.1 esxtop

esxtop provides a real-time view of the ESX server workload. This data can be viewed real-time or redirected to a batch file for further analysis. esxtop usage is out of scope for this document; see Interpreting esxtop 4.1 Statistics for detailed information. The following table summarizes a few of the counters that are useful when monitoring performance for a XenApp workload.

Table 1. esxtop Counter Reference

Display	Metric	Threshold	Description
CPU	%RDY	10	If the threshold is exceeded, over-provisioning of vCPU, excessive usage of vSMP or a limit (check %MLMTD) has been set. This %RDY value is the sum of all vCPUs %RDY for a virtual machine. For example, if the max value of %RDY of 1vCPU is 100% and 4vCPU is 400%. If %RDY is 20 for 1 vCPU then this is problematic, as it means 1 vCPU is waiting 20% of the time for VMkernel to schedule it.
CPU	%CSTP	3	If the threshold is exceeded this indicates excessive usage of vSMP. Decrease amount of vCPUs for this particular virtual machine.
CPU	%MLMTD	0	If larger than 0 the worlds are being throttled. Possible cause is a limit on CPU.
CPU	%SWPWT	5	If the threshold is exceeded the virtual machine is waiting on swapped pages to be read from disk. You may have overcommitted memory.
MEM	MCTLSZ	1	If larger than 0, host is forcing virtual machine to inflate balloon driver to reclaim memory as the host is overcommitted.
MEM	SWCUR	1	If larger than 0 host has swapped memory pages in the past. You may have overcommitted.
MEM	SWR/s	1	If larger than 0 host is actively reading from swap. This is caused by excessive memory overcommitment.
MEM	SWW/s	1	If larger than 0 host is actively writing to swap. This is caused by excessive memory overcommitment.

© 2011 VMware, Inc. All rights reserved.

MEM	N%L	80	If less than 80, virtual machine experiences poor NUMA locality. If a virtual machine has memory size greater than the amount of memory local to each processor, the ESX scheduler does not attempt to use NUMA optimizations for that virtual machine.
NETWORK	%DRPTX	1	If larger than 0 transmit packets are being dropped, hardware is overworked due to high network utilization.
NETWORK	%DRPRX	1	If larger than 0 receive packets are being dropped, hardware is overworked due to high network utilization.
DISK	GAVG	25	Look at DAVG and KAVG as GAVG = DAVG + KAVG.
DISK	DAVG	25	At this level and higher you have disk latency that is likely to be caused by storage array.
DISK	KAVG	2	If 2 or higher disk latency may be caused by the VMkernel. High KAVG usually means queuing. Check QUED.
DISK	QUED	1	If 1 or higher the queue has maxed out. Possibly queue depth is set too low. Check with array vendor for optimal queue value.
DISK	ABRTS/s	1	Aborts issued by virtual machine because storage is not responding. For Windows virtual machines this happens after 60-second default. Can be caused by path failure, or storage array is not accepting IO.
DISK	RESET/s	1	The number of commands resets per second.

5.1.2 vSphere Client

VMware vCenter Server and the vSphere Client are key components when managing a vSphere environment. As a single-pane of glass into the virtual environment, the vSphere Client provides the interface to manage and monitor the virtual infrastructure. vCenter Server provides monitoring, alerting, and real-time and historical graphical performance monitoring. vSphere administrators can use the capabilities of the vSphere Client to establish baseline performance characteristics. These baselines can in-turn be used to adjust built-in alarm thresholds, or create new alarms to meet the monitoring needs of the organization.

5.1.2.1. vCenter Server Alarms

Out of the box, vCenter Server provides over 40 pre-defined alarms that monitor hardware status, resource utilization, IO connectivity, licensing, and more. The built-in alarms can be edited to trigger based on thresholds that match your service-level agreements. By default, the built-in alarms are configured with either no action, or to send a SNMP trap based on the vCenter Server settings. As part of vCenter Server setup, review the built-in alarms and adjust them to fit your requirements, including actions such as sending an email or running a script.

© 2011 VMware, Inc. All rights reserved.

5.1.2.2. vCenter Server Performance Graphs

The performance graphs available within the vSphere Client provide a graphical representation of the performance data collected for objects managed by vCenter Server. This data includes real-time and historical CPU, memory, network, and storage metrics. An administrator can quickly review resource utilization using the overview pane, or switch to the advanced view to drill down into specific component utilization and performance.

The counters that are retained for historical reference are limited by default to minimize space requirements. If it is found that more granular historical data is required the statistics level may be raised from the vCenter Server Settings dialog box. vCenter estimates the space required based on the number of managed objects and the statistics level selected.

Not all counters found in esxtop are captured by vCenter Server. Table 2 lists the esxtop counters listed in Table 1 along with their vSphere client equivalent. Those counters not available through the vSphere Client are also noted.

Table 2. esxtop – vCenter Counter Mapping

Component	esxtop Counter	vSphere Client Counter
CPU	%RDY	Ready
CPU	%CSTP	none
CPU	%MLMTD	none
CPU	%SWPWT	Swap Wait
Memory	MCTLSZ	Balloon
Memory	SWCUR	Swapped
Memory	SWR/s	Swap In Rate
Memory	SWW/s	Swap Out Rate
Memory	N%L	none
Network	%DRPTX	none
Network	%DRPRX	none
Disk	GAVG/cmd	none
Disk	DAVG/cmd	Physical Device command Latency
Disk	KAVG/cmd	Kernel Command Latency
Disk	QUED	none
Disk	ABRTS/s	Command Aborts
Disk	RESET/s	Bus Resets

© 2011 VMware, Inc. All rights reserved.

5.2 Monitoring Citrix

Tools built into the VMware products can provide a granular look into the resource consumption and performance characteristics from an infrastructure point of view. When providing a service such as XenApp virtualized application delivery it is important to also monitor from the point of view of the end-user. Doing this requires tools that understand the applications to be monitored and the activities occurring within the applications being delivered as well as at the XenApp level.

5.2.1 EdgeSight for XenApp

In an existing XenApp environment it is likely that EdgeSight for XenApp is already providing monitoring capabilities. When virtualizing a XenApp farm EdgeSight can continue to provide end-to-end application performance monitoring just as it does in the physical environment.

EdgeSight can provide visibility into the performance of an application, a specific session, and the surrounding infrastructure. When troubleshooting, this can quickly help determine if focus needs to be on the application, XenApp itself, or a component within the infrastructure. Besides monitoring for problems EdgeSight can help optimize the infrastructure by providing real-time and historical performance statistics, as well as application and user-level resource consumption. By monitoring this data virtual machine building blocks that may have been over-provisioned because of lack of real-world data can be adjusted to obtain ideal virtual machine sizing leading to more efficient resource utilization.

© 2011 VMware, Inc. All rights reserved.

6. vSphere Enhancements for Deployment and Operations

You can leverage vSphere to provide significant benefits in a virtualized XenApp environment, including:

- Increased operational flexibility and efficiency – Rapid service deployment in shorter time frames**.**
- Efficient change management – Increased productivity when testing the latest software patches and upgrades.
- Minimized risk and enhanced IT service levels – Zero downtime maintenance capabilities, rapid recovery times for high availability, and compute resource load-balancing.

6.1 Rapid Provisioning and Deployment

Quickly reacting to ever changing requirements allows for a more agile and efficient environment. In a physical environment, reacting to a sudden increase of end-users due to mergers, department consolidations, or acquisitions, can take days or even weeks. vSphere provides the capabilities to turn services around quickly, bringing the time to market down to hours or minutes. By creating a golden template of your XenApp server virtual machine you can be sure that requirement changes can be quickly addressed without the need to provision new hardware. Templates can be deployed as needed and customized as they are deployed using custom scripts, if required. Keeping your templates up to date with the latest security patches results in faster provisioning times because fewer reboots are required after the initial deployment.

6.2 Efficient Change Management

Security patches and software updates are needed more often than ever before. Many new patches and updates are being released, but you have to be sure one of them is not going to bring down your application. In some cases, organizations deploy entire parallel environments to make sure they have the identical configurations in place to test these updates. The problem with these test environments is that they can quickly become outdated. By using VMware snapshots and clones you have more options to test patches before rolling them out to production.

VMware snapshots preserve the state of a virtual machine at a specific point in time, running or powered off. After a snapshot is taken any changes made to the virtual machine are written to a delta file. The delta file includes software, OS patches, and any other changes made within the guest OS of the virtual machine. After validating the update you can choose to commit the changes, or remove the changes and revert to the state at the time of the snapshot.

Cloning a virtual machine may be more appealing to some customers. A clone is an exact replica of your virtual machine and may be taken with the virtual machine powered off (cold-clone) or while the virtual machine is powered on (hot-clone). With an exact replica of the virtual machine you can apply any patches or updates while the clone is disconnected from the network and verify functionality. Validation using clones may be preferable for virtual machines where an application or OS update writes information to other data repositories.

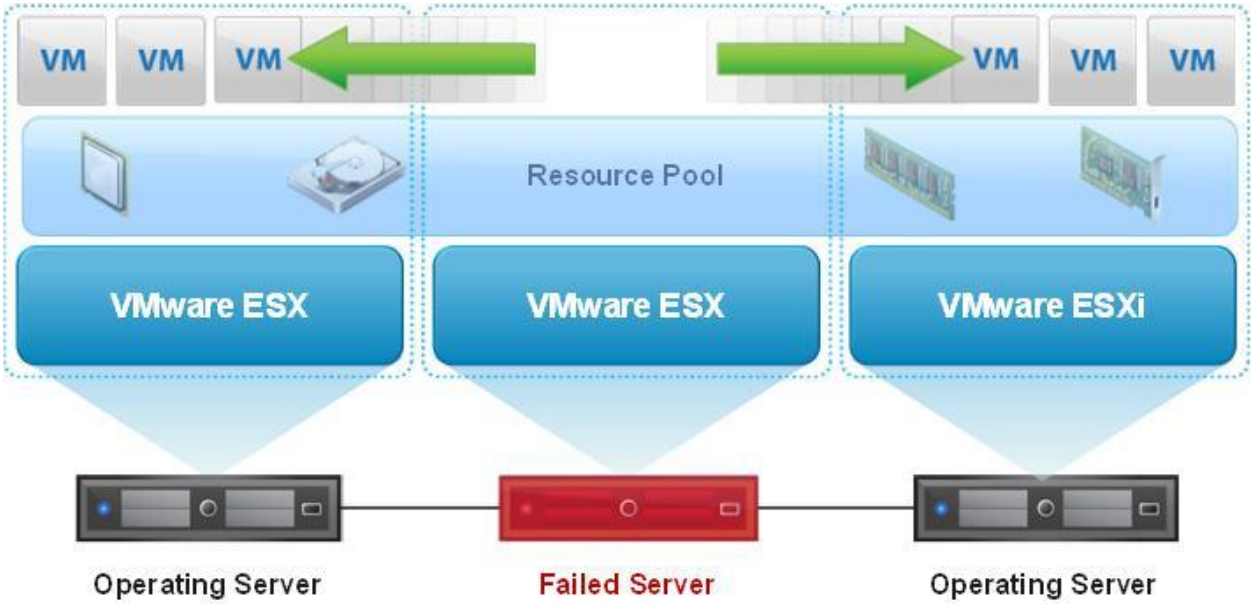
© 2011 VMware, Inc. All rights reserved.

6.3 Enhanced Availability with VMware HA

Application availability is often enabled by implementing complex clustering technologies. These technologies often require advanced licensing, additional hardware, and an advanced skill set to configure and manage. Frequently, these clustering technologies are specific to the protected application.

VMware HA brings high-availability to the infrastructure, removing the complexities from the applications and operating systems. When a vSphere cluster is enabled with VMware HA, cluster nodes maintain heartbeat communication. If failure of a vSphere cluster node is detected virtual machines previously hosted on the failed node are powered up on a surviving host. VMware HA uses a common interface to provide high availability for all applications and operating systems running on vSphere.

Figure 5. VMware HA



If you use VMware HA, be aware that:

- VMware HA handles ESX host hardware failure, but does not monitor the status of the XenApp services—these must be monitored separately.
- VMware HA can monitor the guest operating system using VMware Tools heartbeat and by monitoring virtual machine IO. If a guest OS fails, VMware HA can restart the virtual machine.
- Proper DNS hostname resolution is required for each ESX host in a VMware HA cluster.
- VMware HA heartbeats between hosts are sent via the vSphere VMkernel network, so redundancy in this network is recommended.

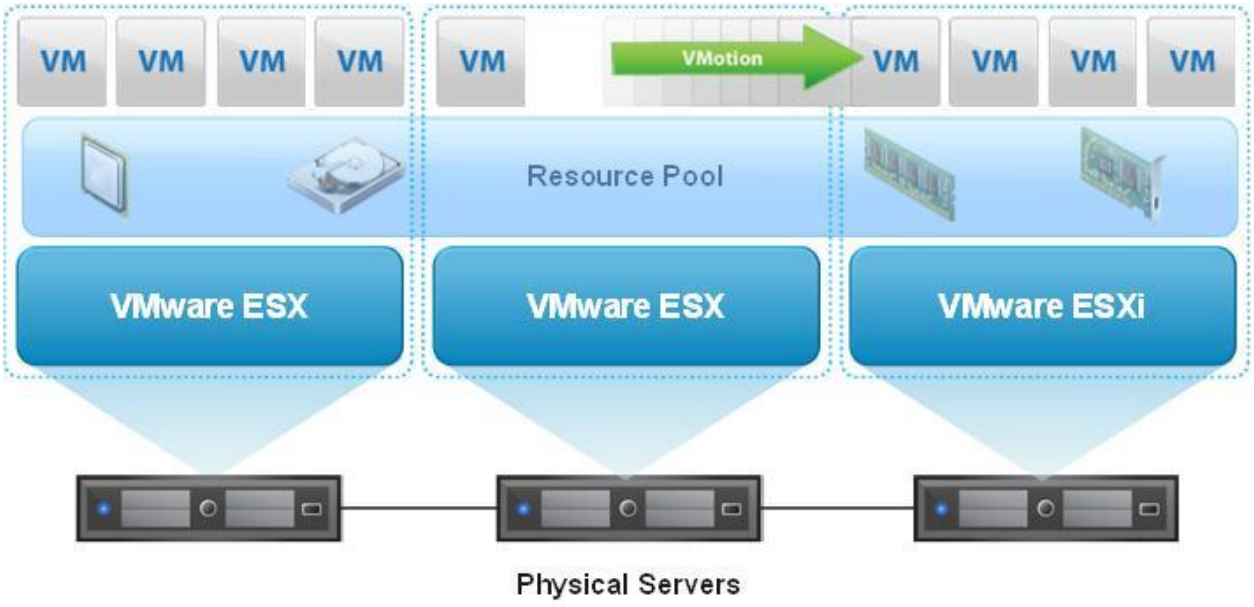
© 2011 VMware, Inc. All rights reserved.

6.4 Resource Optimization with VMware DRS and vMotion

Physical XenApp deployments are bound by the underlying hardware. This often leads to over-provisioning and underutilized hardware. The ability to add and remove compute resources as needed greatly improves the flexibility and efficiency of your XenApp environment. VMware vMotion technology enables the migration of virtual machines from one physical server to another without service interruption. Using vMotion, XenApp virtual machines can be migrated from a heavily-loaded server to one that is lightly loaded. Migrations using vMotion can be initiated on demand, or automated using DRS.

VMware DRS takes the VMware vMotion capability a step further by adding an intelligent resource scheduler. DRS enables you to set resource assignment policies that reflect business needs. VMware DRS does the calculations and automatically handles the details of physical resource assignments. It dynamically monitors the workload of the running virtual machines and the resource utilization of the physical servers within a cluster. When workloads become unbalanced DRS can migrate virtual machines to hosts with additional capacity using vMotion.

Figure 6. VMware DRS and vMotion



VMware vMotion and VMware DRS perform best under the following conditions:

- The source and target ESX hosts must be connected to the same gigabit network and the same shared storage.
- A dedicated gigabit network for VMware vMotion is recommended.
- The destination host must have enough resources.
- The virtual machine must not use physical devices such as CD-ROM or floppy.
- The source and destination hosts must have compatible CPU models, or migration with VMware vMotion fails. For a listing of servers with compatible CPUs refer to VMware vMotion compatibility guides from specific hardware vendors.
- To minimize network traffic it is best to keep virtual machines that communicate with each other together on the same host machine.
- Virtual machines with smaller memory sizes are better candidates for migration than larger ones.

© 2011 VMware, Inc. All rights reserved.

6.5 Security with VMware vShield Endpoint

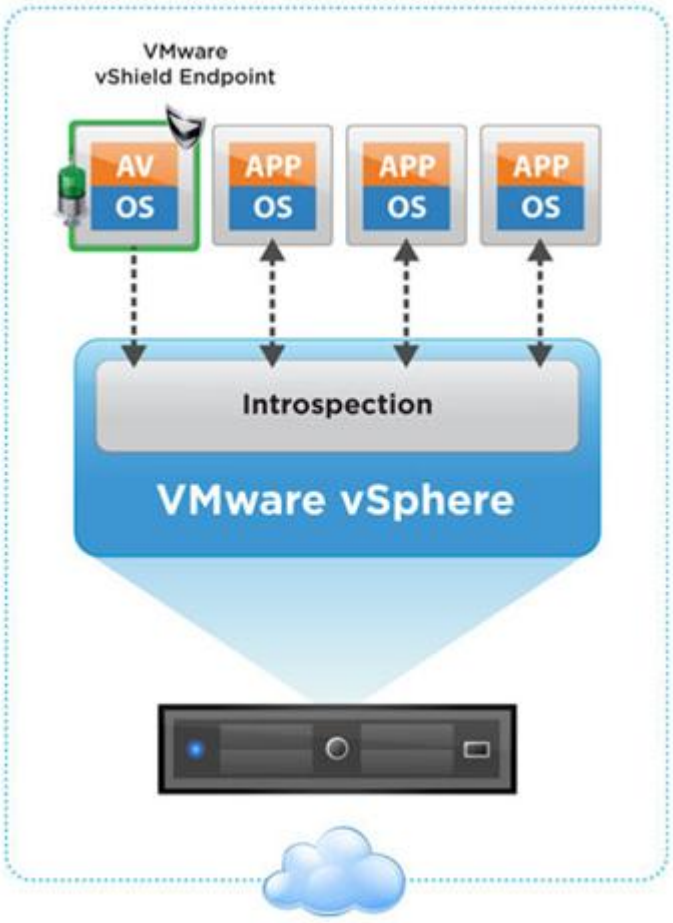
Antivirus protection is critical for production environments. Short of complete isolation with no possibility of external influences there's a pretty good chance that systems will become infected at one time or another. This threat increases dramatically when providing a service where the main interaction comes directly from end users.

Traditional antivirus agents installed locally in the operating system have helped mitigate these threats. On-access file level scanning provides real-time protection for frequently accessed files, while on-demand scanning rounds out the protection by scanning all files on a system. With threats constantly emerging and evolving, updates are critical. The frequency of these updates is different from vendor to vendor, but typically range from hourly to daily. Constant file scanning, network queries for updates and the application of those updates has wreaked havoc on large environments and been the cause of the hourly or daily slowdowns that affect users. Not only does this impact end-user productivity, it can impact perception of service levels, operational efficiency, and overall infrastructure performance.

VMware vShield Endpoint offloads antivirus and anti-malware processing to dedicated security–hardened virtual machines provided by VMware partners. vShield Endpoint consists of the hardened virtual machine, a driver for virtual machines to offload file events, and the VMware Endpoint security (EPSEC) loadable kernel module to link the first two components at the hypervisor layer. vShield Endpoint provides the following benefits:

- Streamline antivirus and anti-malware deployment – Deploy an enterprise antivirus engine and signature file to a single security virtual machine instead of each and every individual virtual machine on a vSphere host.
- Improve virtual machine performance – Securely achieve higher consolidation ratios using the same offload mechanism described in the previous bullet item.
- Prevent antivirus storms and bottlenecks – Prevent antivirus storms and bottlenecks associated with multiple simultaneous antivirus and anti-malware scans and updates.
- Protect antivirus security software from attack – Deploy and run the antivirus and anti-malware client software in a hardened security virtual machine to prevent targeted attacks.

Figure 7. Security with VMware vShield Endpoint



© 2011 VMware, Inc. All rights reserved.